

Thermodynamic Hypothesis as Social Choice: An Impossibility Theorem for Protein Folding

Hammurabi Mendes, Sorin Istrail
{hmendes,sorin}@cs.brown.edu
Brown University

Abstract

Protein Folding is concerned with the reasons and mechanism behind a protein's tertiary structure. The thermodynamic hypothesis of Anfinsen postulates an universal energy function (UEF) characterizing the tertiary structure, defined consistently across proteins, in terms of their aminoacid sequence.

We consider the approach of examining multiple protein structure descriptors in the PDB (Protein Data Bank), and infer individual preferences, biases favoring particular classes of aminoacid interactions in each of them, later aggregating these individual preferences into a global preference. This 2-step process would ideally expose intrinsic biases on classes of aminoacid interactions in the UEF itself. The intuition is that any intrinsic biases in the UEF are expressed within each protein in a specific manner consistent with its specific aminoacid sequence, size, and fold (consistently with Anfinsen's thermodynamic hypothesis), making a 1-step, holistic aggregation less desirable.

Our intention is to illustrate how some impossibility results from voting theory would apply in this setting, being possibly applicable to other protein folding problems as well. We consider concepts and results from voting theory and unveil methodological difficulties for the approach mentioned above. With our observations, we intend to highlight how key theoretical barriers, already exposed by economists, can be relevant for the development of new methods, new algorithms, for problems related to protein folding.

1 Introduction

The systematic aggregation of individual preferences into a global social preference is the principal concern of social choice theory. Its origins trace back to the 13th-century, when the philosopher Ramon Lull considered a voting method that counted pairwise preferences among candidates. This technique was later recovered by the famous 18th-century mathematicians Borda and Condorcet. Later, in 1951, Kenneth Arrow stated his famous impossibility theorem ([3], for the version of 1963), which is viewed as a landmark of social choice theory. This important result defines apparently innocuous properties that certain preference aggregation functions (called social welfare functions) should respect, but also demonstrates that such properties are conflicting.

Many other impossibility results followed. The conclusions usually result from an axiomatic setting, where aggregation functions are modeled after axioms defining their intended behavior. Also common is the relative unconcern for the elements subject to preferences. Many aggregation settings, in addition, do **not** consider voters as players, but merely as providers of information – and we are interested on those settings **only**. Note that the observations above suggest a potential application of the impossibility results to other circumstances that involve aggregation. In this work, we model a problem related to protein folding as an aggregation function prone to those impossibility results, exploring the connection between these ideas.

The protein folding problem is basically interested on how and why proteins settle in particular three-dimensional structures. The well-accepted *thermodynamic hypothesis* of Anfinsen [2] states that the protein structure, informally speaking, is a unique, stable, and viable state depending solely on *aminoacid sequence*. The state minimizes an energy function accounting for physiochemical utilities. Such *universal energy function* (UEF) is defined consistently across proteins, in terms of the particular aminoacid interactions that define each of them.

The protein folding literature often takes a *holistic* approach to analyzing the PDB [1] in order to obtain certain conclusions from stable folds. We define a problem related to protein folding where a 2-step aggregation process will seem more intuitive: suppose we intend to infer hypothesized intrinsic preferences on the UEF, which would favor certain classes of aminoacid interactions over others. The idea is to first aggregate multiple *individual preferences*, and later consolidate them into a single *global preference*. In other words, we consider and explore the following scenario:

- Say that the UEF has an intrinsic *global preference* among interaction classes. For instance, the UEF might generally favor (A, V) , the interaction of the aminoacids alanine and valine, over (G, L) , the interaction of aminoacids glycine and leucine.
- This global preference is manifested distinctly across proteins, in light of the thermodynamic hypothesis. In other words, a protein’s stable fold implies an *individual preference* among interaction classes, which results from its corresponding aminoacid sequence and the global preference intrinsic to the UEF.

We consider the possibility of inferring or estimating the postulated *global preference* among interaction classes, intrinsic to the UEF, by aggregating *individual preferences*, extracted separately from multiple stable folds in the PDB. We assess and qualify aminoacid interactions inside each protein in order to generate the individual preferences. Importantly, we assume that aminoacid interactions are totally defined by *pairwise interactions* (see Sec. 2.2). The existence and expression of a pairwise interaction depends on the aminoacids involved (their physiochemical affinities) and their spacial arrangement (their distance and contact).

We particularly think that such 2-step aggregation process is particularly sensible in light of Anfinsen’s thermodynamic hypothesis. The idea is that individual proteins may express universal biases differently, according to their particular aminoacid sequence, size, and fold (note here the thermodynamic hypothesis). If we holistically analyze the PDB, performing instead a 1-step aggregation process, we effectively aggregate over aminoacid interactions from multiple proteins. By assumption, those aminoacid interactions may express biases differently. This argument might as well be applicable to other problems related to protein folding where the thermodynamic hypothesis is considered.

The problem formalized before is formally *weaker* than those aiming for more complete characterizations of the UEF with our 2-step aggregation process and the same axiomatic setting. Here, while we do not intend to completely characterize the UEF, we do intend to identify intrinsic biases among interaction classes in the UEF. Unfortunately, this weaker problem will still manifest an impossibility in light of some seemingly reasonable assumptions (discussed in Sec. 2 and Sec. 4). Stronger characterizations of the UEF via the method above will suffer from similar methodological limitations unless axioms are relaxed, made more specific, or otherwise changed appropriately.

Our intention and contribution is to highlight theoretical impossibility results for the problem considered above, under reasonable assumptions (see Sec. 4). We use concepts and results from voting theory, a distinct investigation field of social choice theory, in Economics. The goal is to sparkle interest in the profound results from voting theory, and highlight a connection of this area with methods that, so far, have been mainly a concern of biologists, chemists, physicists, and computer scientists. We note that the remarks made here result from a suitable axiomatic framework, from the application of voting theory concepts, and from direct reduction to previous results in this area. We intend to suggest how some key theoretical barriers, already exposed by economists, can be relevant for the development of new methods, new algorithms, for problems related to protein folding, particularly those in which the thermodynamic hypothesis is considered.

2 The Analytic Framework

In this section, we present the axioms and definitions that form our analytic framework. We adhere to the thermodynamic principle of [2], and further assume that aminoacid interactions in the universal energy function are totally defined by pairwise interactions. We try to abstract concepts of purely biological nature as much as possible. Again, the purpose is to bring impossibility results from voting theory to our 2-step preference model and suggest that they might be relevant to other problems related to protein folding.

2.1 Thermodynamic Hypothesis

There are two principal questions in the Protein Folding research [9]. The first one concerns the reasons why proteins settle in a specific tertiary structure. The second asks how to correctly predict the tertiary structure given the protein’s aminoacid configuration.

The thermodynamic hypothesis [2] says that the tertiary structure results solely from *aminoacid interactions*, maximizing a function of utility character and physiochemical nature. Also, the final structure is (1) *unique*, as other possible states have strictly lower energy; (2) *stable*, as small external changes do not disrupt the structure; and (3) *viable*, smoothly accessible by physiochemical changes.

Axiom 2.1. *The thermodynamic hypothesis of [2] is valid, hence the protein’s tertiary structure is unique, stable, viable, and depends solely on aminoacid interactions.*

2.2 Pairwise Interactions

Previous literature suggests energy functions depending on number, distance, and additional physiochemical properties (hydrophobia/hydrophilia, among others) between *pairs* of aminoacids in the three-dimensional protein structure [11, 17, 13]. We will follow such approach, hence *pairwise interactions* fully characterize our aminoacid interactions in this work.

Identifying pairwise interactions instead of trying to obtain all collective, grouped interactions is much more viable computationally. Besides, it is conceivable that more intricate interaction patterns be extrapolated from pairwise interactions. In this work, we consider pairwise energy functions, i.e., those totally defined by energy potentials on pairwise interactions among aminoacids.

Axiom 2.2 (Pairwise Hypothesis). *All existent information on aminoacid interactions can be obtained by considering only pairwise interactions.*

2.2.1 Interaction Classes and Instances

An *interaction class* is a pair (X_1, X_2) where X_1 and X_2 are *names* of aminoacids. For instance, (A, V) is the interaction of the aminoacids alanine and valine, and (G, L) is the interaction of aminoacids glycine and leucine. We consider (X_1, X_2) equal to (X_2, X_1) for any aminoacid names X_1, X_2 . Considering the standard 20 aminoacids that characterize proteins, we have 380 possible interactions classes.

Within a protein’s stable fold, an *interaction instance* of (X_1, X_2) *occurs* if the distance between two aminoacids of types X_1 and X_2 , is within an arbitrary threshold τ . In practice, the distance could be taken over heavy atoms, centroids in aminoacids, or c_α locations in aminoacids (the location of the central carbon atom). Our results are independent of the choice above. We represent interaction classes with uppercase letters (e.g. I), and interaction instances with lowercase letters (e.g. i). In a protein’s stable fold P , the interaction instances of the interaction class I defines the set $Inst_P(I)$.

Establishing a distance threshold might seem arbitrary and even unrealistic, as overlooking interaction instances that barely fail such distance may discard useful information. Also, whatever account attributed to interaction instances must also consider certain physiochemical properties, for instance, hydrophobia or hydrophilia¹, usually compiled into tables such as the Miyazawa-Jernigan [13]. One could argue that the right choice of distance threshold and tables describing physiochemical interactivity could be enough to totally define the UEF. Let us fix (yet not specify) τ big enough as to accommodate all possible *actual* aminoacid interactions, and *assume* that contacts are scored accordingly.

Therefore, to each interaction instance i , we associate an *energetic score* $e(i) \in \mathbb{R}$, which abstracts the following parameters: (1) distance threshold τ ; (2) contact point; (3) the physiochemical properties of the aminoacids involved. As far as identifying intrinsic preferences over interaction classes in the UEF, our results are independent of the above choices – in other words, independent of the energetic score function. The key to such generality of results lies in our 2-step preference model, in our axiomatic setting, and impossibility results from voting theory.

¹Respectively, the tendency or aversion to bind to water.

3 Internal Aggregation

The *internal aggregation* is the procedure where the energetic scores of the interaction instances are taken into account and define a protein’s individual preference. For any protein P :

- Each interaction instance i has an energetic score $e(i) \in \mathbb{R}$ modeling distance threshold, contact point, and physiochemical properties of i . The function $e()$ is fixed yet arbitrary. The value of $e(i)$ can be positive or negative, modeling physiochemical properties that might manifest with opposing effect (for instance, hydrophobia vs. hydrophilia).
- Each interaction class I has a set of interaction instances $Inst_P(I)$. The individual preference of P regarding I_1 versus I_2 depends *only* on $Inst_P(I_1)$ and $Inst_P(I_2)$, and the energetic scores respectively associated with each interaction instance. We consider two scenarios:
 1. For each interaction instance I , we can infer an utility value $u_P(I) \in \mathbb{R}$, completely defined in terms of $Inst_P(I)$. Considering the same protein/function, the utilities are comparable: if $u_P(I_1) > u_P(I_2)$, then P manifests preference (i.e., bias) for I_1 instead of I_2 ; if $u_P(I_1) = u_P(I_2)$, then P does not manifest any preference (i.e., bias) for I_1 instead of I_2 .

Note that we do *not* assume comparability between utility levels across two different proteins. This is what characterizes our 2-step preference model: utilities are meaningful to its corresponding protein only. In Sec. 4.4 we in fact vary the comparability between utility levels across proteins, and explore the consequences of each mode.

2. We can only infer a total order \preceq_P among interaction classes, completely defined in terms of $Inst_P(I_1)$ and $Inst_P(I_2)$. If $I_1 \succ_P I_2$, then P prefers I_1 to I_2 ; if $I_1 \prec_P I_2$, then P prefers I_2 to I_1 ; if $I_1 =_P I_2$, then P is indifferent regarding I_1 versus I_2 .

We call the first scenario the *utility* preference model, and the second scenario is an *ordinal* preference model (our terminology). We examine the two scenarios separately, in each case we apply the appropriate concepts and impossibility results from social choice theory. Recapitulating, our goal is to evaluate the viability of the following strategy:

Internal Aggregation: Examine interaction instances of each protein, deriving the manifested *individual preferences*, either a utility function u_P or a total order \preceq_P , for each protein P .

External Aggregation: Examine individual preferences and derive the *global preference*, which is similarly a utility function or a total order.

We assume that the internal aggregation might be subject to errors, but not in presence or scope such that they might deem any individual preference “more representative” than another. Errors are small and covert. This idea is formalized by the anonymity requirement (Axiom 4.5) for the external aggregation.

4 External Aggregation

In this section we will axiomatize the requirements of the external aggregation procedure. We discuss the general problem structure on Sec. 4.1, followed by an axiomatic formulation of the basic requirements on Sec. 4.2. We finally show that the external aggregation (axiomatized as below) is

subject to fundamental impossibility results studied in the context of voting theory regardless of: (1) the function e ; and (2) the utility function u_P or the total order \preceq_P for any protein P .

In other words, the strategy outlined above, in the end of Sec. 3, will pose some fundamental methodological limitations *regardless* on how the internal aggregation is defined – how distances, contact, physiochemical potentials, as well how preferences are extracted. Specific axioms and impossibility theorems for the case with ordinal preference is presented on Sec. 4.3, and for the case with utility preference on Sec. 4.4.

4.1 Basic Notation

The external aggregation is defined as follows. Consider all interaction classes $\mathcal{I} = \{I_1, \dots, I_{380}\}$, and arbitrary n proteins $\mathcal{P} = \{P_1, \dots, P_n\}$, with some finite $n \geq 2$. Considering our two preference models:

- In the ordinal preference setting, the individual preference of P_i is a *total order* \preceq_i over \mathcal{I} . By definition, these relations are total ($I_a \preceq_i I_b$ or $I_b \preceq_i I_a$ for any I_a, I_b), transitive (if $I_a \preceq_i I_b$ and $I_b \preceq_i I_c$ then $I_a \preceq_i I_c$), and antisymmetric ($I_a \preceq_i I_b$ and $I_b \preceq_i I_a$ implies $I_a = I_b$).
- In the utility preference setting, the individual preference of P_i is a *utility function* $u_i : \mathcal{I} \rightarrow \mathbb{R}$. The utilities imply a total order (if $u_i(I_a) \leq u_i(I_b)$ then $I_a \preceq_i I_b$).

The *individual preference* manifested by protein P_i is denoted as $\mathbb{I}(i)$. The *preference profile* is composed of all individual preferences: $\mathbb{I} = (\mathbb{I}(0), \dots, \mathbb{I}(n))$. Finally, the *universe of preference profiles* is composed of all preference profiles: $\mathbb{I}^* = \{\mathbb{I} : \mathbb{I} \text{ is a valid preference profile}\}$. The external aggregation procedure returns a *global preference*, either a total order $\preceq_{\mathbb{O}}$ or a utility function $u_{\mathbb{O}}$ over \mathcal{I} . We consider the two cases separately.

In terms of notation, for any individual or the global preference X , if $I_a \preceq I_b$ in X , we write $I_a \preceq_X I_b$ (and similarly for other relations). So, if $I_a \preceq I_b$ in $\mathbb{I}(i)$, we write $I_a \preceq_{\mathbb{I}(i)} I_b$; besides, if $I_a \preceq I_b$ in \mathbb{O} , we write $I_a \preceq_{\mathbb{O}} I_b$. For any two preferences X and Y , if $I_a \preceq_X I_b \Leftrightarrow I_a \preceq_Y I_b$, then we write **SameOrder**($X, Y, I_a \sim I_b$).

4.2 General Axioms

Our first axioms simply models the requirement that the procedure receives multiple individual preferences and produces a single global preference. Individual preferences are transitive ordering over interaction classes.

Axiom 4.1. Agreement *After the external aggregation, we obtain a single global preference among interaction classes.*

Axiom 4.2. Transitivity *The individual preferences as well as the global preference are total transitive orderings over interaction classes.*

The unrestricted domain requirement, stated below, is an imposition on the external aggregation, not on the actual preference profile. Intuitively speaking, the external aggregation must be able to handle arbitrary transitive individual preferences.

Axiom 4.3. Unrestricted Domain *Any total order among interaction classes is permitted in a preference profile.*

If any set of at least three interaction classes $S = \{I_1, I_2, I_3\}$ is such that both (1) $I_i \succeq I_j$ and (2) $I_j \succeq I_i$ are possible preferences in the preference profile, for any $1 \leq i, j \leq 3$, S is called a 3-unrestricted preference set.

In a first glance, Axiom 4.3 may appear overly imposing for practical purposes. After all, the preference profile expressed in the PDB, given a suitable internal aggregation procedure, can be more restrictive. However, it seems reasonable to expect a 3-unrestricted preference set in the preference profile after the internal aggregation. In that case, the Arrow’s Impossibility Theorem (Theorem 4.14) is applicable, and May’s Theorem (Theorem 4.15) may imply the possibility of breaking transitivity (discussed ahead). Even considering the much less restrictive condition of having only a 3-unrestricted preference set in the preference profile, the results apply. We keep the complete unrestricted domain axiom for its generality appeal.

The unanimity requirement requires that if all proteins expose a preference on a particular interaction class over another, the unanimity should be respected.

Axiom 4.4. *Unanimity.* For any $\mathbb{I} \in \mathbb{I}^*$, if **SameOrder**($\mathbb{I}(i), \mathbb{I}(j), I_a \sim I_b$) for all P_i, P_j then **SameOrder**($\mathbb{I}(i), \mathbb{O}, I_a \sim I_b$) for all P_i , considering arbitrary I_a, I_b ;

We assume that proteins manifest preferences over interaction classes in an indistinguishable manner; i.e., we cannot infer whether protein A manifests a “better” individual preference than protein B . As mentioned before, we assume that the internal aggregation might be subject to errors, but not in presence or scope such that they might deem any individual preference “more representative” than another. Errors are small and covert. In specific terms, the preferences manifested by protein A and those by protein B are considered without regard to their identity.

Axiom 4.5. *Anonymity.* For any $\mathbb{I}', \mathbb{I}'' \in \mathbb{I}^*$, whenever $\mathbb{I}' = (\mathbb{I}(0)', \dots, \mathbb{I}(n)')$ is a permutation of $\mathbb{I}'' = (\mathbb{I}(0''), \dots, \mathbb{I}(n''))$, we have that $\mathbb{O}' = \mathbb{O}''$.

We note that anonymity implicates in non-dictatorship, i.e., the fact that no individual preference completely and exclusively defines the global preference.

Axiom 4.6. *Non-Dictatorship.* For all $P_i \in \mathcal{P}$ there exists at least one $\mathbb{I} \in \mathbb{I}^*$ and elements I_a, I_b such that $I_a \succ_{\mathbb{I}(i)} I_b$ but $I_a \preceq_{\mathbb{O}} I_b$.

4.3 Ordinal Preference Model

Our external aggregation procedure is modeled after *social welfare functions*, a basic concept in social choice theory, simply representing a procedure that receives total orders on elements of \mathcal{I} and produces a total order over elements of \mathcal{I} . We start with basic axioms meaningful for the ordinal preference model, followed by impossibility results.

4.3.1 Axioms

The neutrality requirement, formalized below, basically requires that all pairwise preferences over interaction classes are treated equally: the procedure computing the outcome of all pairwise preferences is oblivious to the particular names of interaction classes. We say that two preference profiles are *compatible* whenever they have the same number of individual preferences.

Axiom 4.7. *Neutrality.* For any compatible $(\mathbb{I}', \mathbb{I}'')$ and corresponding outputs $(\mathbb{O}', \mathbb{O}'')$, if we have $[\forall P_i e_a \succeq_{\mathbb{I}(i)'} e_b \Leftrightarrow e_c \succeq_{\mathbb{I}(i)''} e_d]$ then $e_a \succeq_{\mathbb{O}'} e_b \Leftrightarrow e_c \succeq_{\mathbb{O}''} e_d$, considering arbitrary e_a, e_b, e_c, e_d ;

If we make $e_c = e_a$ and $e_d = e_b$, neutrality implies the famous and controversial Independence of Irrelevant Alternatives (IIA). This requirement basically says that for any two aggregation scenarios where individual preferences over arbitrary elements e_a and e_b remain identical, the algorithm output similarly remains identical over those elements.

Axiom 4.8. Independence of Irrelevant Alternatives (IIA). *For any compatible $\mathbb{I}', \mathbb{I}'' \in \mathbb{I}^*$ and corresponding outputs $\mathbb{O}', \mathbb{O}''$, if we have $[\forall P_i \text{ SameOrder}(\mathbb{I}(i)', \mathbb{I}(i''), e_a \sim e_b)]$ then we have $\text{SameOrder}(\mathbb{O}', \mathbb{O}'', e_a \sim e_b)$, considering arbitrary e_a, e_b ;*

Intuitively speaking, we want an aggregation procedure that behaves similarly to continuous functions in the following sense: small changes in the preference profile will not induce substantial changes in the global preference. We formalize this requirement using distance metrics. We first define a distance metric $d : \mathbb{I}^* \times \mathbb{I}^* \rightarrow \mathbb{R}$, which naturally satisfies the following:

1. $d(X, X) = 0$
2. $d(X, X') = d(X', X)$
3. $d(X, X'') \leq d(X, X') + d(X', X'')$

for any $X, X' \in \mathbb{I}^*$. Then, between two compatible preference profiles $\mathbb{I}', \mathbb{I}''$, we define

$$D = \sum_{i \in \{1, \dots, n\}} d(\mathbb{I}(i), \mathbb{I}(i')).$$

Technically, we want smaller changes in the preference profile to produce smaller changes in the global preference, where changes are measured according to the distance function. The motivation is having an external aggregation that is *robust* to small imprecision in the internal aggregation. Formally:

Axiom 4.9. Proximity Preservation. *There exists a distance function d over \mathbb{I}^* such that*

$$D(\mathbb{I}, \mathbb{I}') \leq D(\mathbb{I}, \mathbb{I}'') \Rightarrow d(\mathbb{O}, \mathbb{O}') \leq d(\mathbb{O}, \mathbb{O}'').$$

We want an external aggregation that responds positively (or at least monotonically) to the preferences of the proteins. We want that positive changes on the preference of I_1 over I_2 , say, due to improving our internal aggregation, induce positive (or at least monotonic) changes in the global order. The two axioms are presented below.

Axiom 4.10. Positive Responsiveness. *For any $\mathbb{I}', \mathbb{I}'' \in \mathbb{I}^*$ and corresponding outputs $\mathbb{O}', \mathbb{O}''$, if $\text{SameOrder}(\mathbb{I}(i)', \mathbb{I}(i''), e_a \sim e_b)$ for all $i \neq j$, and, for some j*

1. $e_a \prec_{\mathbb{I}(j)'} e_b \Rightarrow e_a \succeq_{\mathbb{I}(j)''} e_b$;
2. $e_a =_{\mathbb{I}(j)'} e_b \Rightarrow e_a \succ_{\mathbb{I}(j)''} e_b$,

then $e_a \succeq_{\mathbb{O}'} e_b \Rightarrow e_a \succ_{\mathbb{O}''} e_b$, considering arbitrary e_a, e_b .

Axiom 4.11. Monotonic Responsiveness. *For any $\mathbb{I}', \mathbb{I}'' \in \mathbb{I}^*$ and corresponding outputs $\mathbb{O}', \mathbb{O}''$, if $\text{SameOrder}(\mathbb{I}(i)', \mathbb{I}(i''), e_a \sim e_b)$ for all $i \neq j$, and, for some j*

1. $e_a \prec_{\mathbb{I}(j)'} e_b \Rightarrow e_a \succeq_{\mathbb{I}(j)''} e_b$;
2. $e_a =_{\mathbb{I}(j)'} e_b \Rightarrow e_a \succ_{\mathbb{I}(j)''} e_b$,

then $e_a \succeq_{\mathbb{O}'} e_b \Rightarrow e_a \succeq_{\mathbb{O}''} e_b$, considering arbitrary e_a, e_b .

4.3.2 Impossibility Results

In this section, we romp through impossibility results from social choice theory and expose contradictions hidden in our previous axioms. First, considering the results from [5], we conclude that any external aggregation procedure that respects unanimity and satisfies anonymity does not preserve proximity.

Theorem 4.12. *No external aggregation can simultaneously respect Axioms 4.4, 4.5, 4.9.*

Proof. Proof in [5]. □

Theorem 4.12 exposes a methodological difficulty: an external aggregation procedure that respects anonymity and unanimity does not preserve proximity, which goes against the intuition that smaller changes in the preference profile should imply in smaller changes in the global preference. Note that the theorem in [5] is valid even when the unrestricted domain axiom is not considered. However, if we waive proximity preservation (hence keeping unanimity and anonymity), having an unrestricted domain still frustrates our requirements, as we conflict with neutrality:

Corollary 4.13. *No external aggregation can simultaneously respect Axioms 4.2, 4.3, 4.4, 4.5, 4.7.*

Proof. Since anonymity implicates in non-dictatorship and neutrality implicates in the IIA, this is a corollary of the Arrow's Impossibility Theorem [3]. □

The famous Arrow's Impossibility Theorem [3] is widely considered as a landmark of modern social choice theory. There exist many proofs for such result (including [3, 6, 10]). The original theorem is presented below. Note that even relaxing anonymity to non-dictatorship, and neutrality to the IIA, the original theorem applies:

Theorem 4.14 (Arrow's Impossibility Theorem). *No external aggregation (modeled after social welfare functions) can simultaneously respect Axioms 4.2, 4.3, 4.4, 4.6, 4.8.*

Proof. Proof in [3]. □

If we want a positively responsive aggregation procedure (i.e., respecting axiom 4.10), May's Theorem [12] says that the preference regarding *each* pairwise comparison among interaction classes is necessarily decided by simple majority. For any interaction classes $I_a, I_b \in \mathcal{I}$, define

$$N(I_a > I_b) = |\{P_i : I_a \succ_{\mathbb{I}(i)} I_b\}|.$$

Theorem 4.15 (May's Theorem). *For each interaction class (I_a, I_b) , any decision function² respecting Axioms 4.4, 4.7, and 4.10 is such that $I_a \succeq_{\mathbb{O}} I_b \Leftrightarrow [N(I_a > I_b) \geq N(I_b > I_a)]$.*

The problem implied by May's Theorem is that the global preference might be *intransitive* [6, 4] if a 3-unrestricted preference set exists in the preference profile (which seems reasonable after the internal aggregation). In fact, even waiving neutrality (preserving just the IIA) conflicts with positive responsiveness ([4], chap. 22), so perhaps the problem is the positive responsiveness itself. Settling for monotonicity (Axiom 4.11) works around May's theorem, but then we could argue whether waiving neutrality (preserving just the IIA) is reasonable or not.

²For us, external aggregation procedure.

At first, it sounds reasonable to treat preferences consistently across different interaction classes (as suggested by neutrality and by the IIA). As implied above, restrictions like neutrality (and its weaker restriction, the IIA), even with seemingly natural meaning, have strong axiomatic implications. Nevertheless, dropping neutrality poses serious methodological difficulties for our strategy: in that case, *names* of interaction classes matter for our external aggregation. Their relationship and place, however, are precisely our unknowns.

Many impossibility theorems in the voting theory literature presume unrestricted domains (or simply a 3-unrestricted preference set). For the Arrow’s Theorem as well as the May’s Theorem, performing suitable *domain restrictions*, where we limit the preference profile, might be enough to overcome the theoretical impossibilities. If preferences are *single-peaked* or *quasi-transitive* (please refer to [16, 15]), we also dodge from the impossibility results above. We invite the reader to appreciate these previous results and understand the relevance of these theorems for problems related to protein folding, and perhaps in other similar contexts where a 2-step preference model is similarly suitable.

4.4 Utility Preference Model

In this section, we discuss the utility preference model. In this setting, it is possible to obtain utility values associated with each interaction class in each individual protein. In other words, for any protein P_i , there exists a utility function u_i mapping interaction classes to utility values. Recall that we intend to analyze preferences manifested by individual proteins in an individual setting, and *then* make an external aggregation. We presume utility functions that are not comparable across proteins. In the following subsection, we improve the comparability and explore the consequences.

For now, we technically assume that the utility function for P differ in origin and scale from the utility function for Q , taking any proteins $P \neq Q$. Formally:

Axiom 4.16. *Any utility function is a strictly increasing linear transformation of another utility function: $u_P = \alpha_{pq}u_Q + \beta_{pq}$ for some $\alpha_{pq} > 0$ and any β_{pq} , considering any proteins P and Q .*

Borrowing terminology from voting theory, there is no interpersonal comparisons of utility levels or of gains and losses among different individual preferences. Note that this is fundamentally different from the approach where there is free comparison among utility levels (energy scores) measured across different proteins, as seen in [11, 17, 13].

With that in mind, we are still going to consider ordinal preferences, however expressed through utility functions, in a continuous manner. Consider the space \mathbb{R}_+^{380} where the i -th coordinate of an arbitrary point p is represented by $p[i]$. A point p an *expression descriptor* for interaction classes: p is such that the bigger the $p[i]$, the bigger the associated expression for the interaction class I_i .

Furthermore, for sufficiently close levels of expression descriptors (i.e., points in \mathbb{R}_+^{380}), the utility function has sufficiently close values. Formally, we define

$$U_P(p) = \sum_{I_1 \dots I_{380}} u_P(I_i) \cdot p[i],$$

representing the utility associated with some point p . Technically, U_P is a utility function over the expression descriptor space. Note that U_P is linear, implying that the expression of I_1 is worth $u_P(I_1)/u_P(I_2)$ times the expression of I_2 for any interaction classes I_1 and I_2 and protein P .

Of course, this informational framework is subject to the same impossibility results as before, but other interesting results further apply. We now briefly overview the approach of [7, 14], describing an impossibility result that mirrors (in the continuous setting) Theorem 4.12.

The indifference hyperplane H_P^c is the hyperplane containing points p where $U_P(p) = c$. Note that H_P^c and H_P^d are parallel for any positive constants c and d . The unit vector perpendicular to H_P^c (for an arbitrary constant $c > 0$), called v_P , represents the *direction* where the utility *increases* for P . The length of the vector is disconsidered precisely because we are taking a purely ordinal interpretation of the utility functions: their origin and scale are not comparable among different proteins.

The individual preference for P , denoted by v_P , is a unit vector indicating the direction that maximizes U_P . Note that v_P denotes a unique *point* in the 380-dimensional sphere S^{380} . The global preference, denoted by $v_{\mathbb{O}}$, and is also a unit vector indicating the direction that maximizes the global utility. The external aggregation is such that

$$E : \underbrace{S^{380} \times \dots \times S^{380}}_n \rightarrow S^{380},$$

where E satisfies the following:

Axiom 4.17. *Continuity.* *The external aggregation E is a continuous function.*

The anonymity and unanimity, central requirements for our aggregate mechanism, are rephrased for our continuous setting below.

Axiom 4.18. *Continuous unanimity.* *If $v_P = v$ for all proteins P , then the $v_{\mathbb{O}} = v$.*

Axiom 4.19. *Continuous anonymity.* *$E(v_1, \dots, v_n) = E(\pi(v_1), \dots, \pi(v_n))$, where π is a permutation function over individual preferences.*

Theorem 4.20. *No external aggregation can simultaneously respect Axioms 4.18, 4.19, and 4.17.*

Proof. Follows from [7, 14]. □

Note how Theorem 4.20 mirrors Theorem 4.12 in the continuous setting. At first glance, it seems that anonymity and unanimity are inherently incompatible with a notion of continuity. In fact, the *generality* imposed by seemingly innocuous axioms is also fundamental to such impossibilities. Specifically, performing *domain restrictions*, that is, reducing the generality of possible individual preferences, is often sufficient to escape impossibility results (please refer to [7, 14]).

The continuity requirement models our robustness to imprecision in the internal aggregation procedure. It parallels the proximity preservation property defined for ordinal preferences. In both scenarios, in light of our general axioms, we have impossibility results that seriously hinder the methodological strategy considered in this work.

4.4.1 Inter-Protein Comparison of Energy Levels

Let us now increase the comparability among utility function for different proteins. The utility function for P differs in origin, but not in scale, from the utility function for Q , taking any proteins $P \neq Q$. Formally:

Axiom 4.21. Any utility function is a strictly increasing linear transformation of another utility function: $u_P = \alpha u_Q + \beta_{PQ}$ for some fixed and unique $\alpha > 0$ and any β_{PQ} , considering any proteins P and Q .

Borrowing terminology from voting theory, there is no interpersonal comparisons of utility levels, but there is comparison between gains and losses among different individual preferences, since the scale is the same. For the utility preference model, we write **SameValue**($X, Y, I_a \sim I_b$) to denote $u_X(I_a) = u_Y(I_b)$. Now consider the axioms below.

Utility IIA. For any compatible $\mathbb{I}', \mathbb{I}'' \in \mathbb{I}^*$ and $\mathbb{O}', \mathbb{O}''$, if $[\forall P_i \text{ SameValue}(\mathbb{I}(i)', \mathbb{I}(i''), e_a \sim e_b)]$ then we have that **SameOrder**($\mathbb{O}', \mathbb{O}'', e_a \sim e_b$), considering arbitrary e_a, e_b ;

Strict Unanimity. For any $\mathbb{I} \in \mathbb{I}^*$, if **SameOrder**($\mathbb{I}(i), \mathbb{I}(j), I_a \sim I_b$) for all P_i, P_j then we have **SameOrder**($\mathbb{I}(i), \mathbb{O}, I_a \sim I_b$) for all P_i , considering arbitrary I_a, I_b ; Also, if $I_a \succ_{\mathbb{I}(i)} I_b$ for an arbitrary protein P , then $I_a \succ_{\mathbb{O}} I_b$.

Interestingly, those axioms imply that the external aggregation is the *utilitarian* rule [8]:

$$I_a \succeq_{\mathbb{O}} I_b \Leftrightarrow \sum_{P \in \mathcal{P}} u_P(I_a) \geq \sum_{P \in \mathcal{P}} u_P(I_b).$$

In current literature [11, 17, 13], it is usual to perform a 1-step holistic analysis of the PDB in order to infer a certain conclusion. Social choice theory can provide interesting characterizations and consequences for different levels of comparability of preferences, or whichever concept maps suitably (e.g energy scores).

5 Discussion and Final Remarks

In this work we hope to have illustrated the application and relevance of social choice theory not only to the problem and methodology that intends to infer biases in the UEF, but to other problems related to protein folding where the 2-step approach makes sense. We particularly think that the 2-step approach is sensible in light of Anfinsen’s thermodynamic hypothesis, as if stable folds are the result solely of a protein’s individual aminoacid sequence, considering the PDB holistically may not be always desirable.

Note that the problem considered in this work is formally *weaker* than those aiming for more complete characterizations of the UEF with our 2-step approach and the same axiomatic setting. Stronger characterizations of the UEF via the method above will suffer from similar methodological limitations unless axioms are relaxed, made more specific, or otherwise changed appropriately.

Many impossibility results highlighted here rely on the *generality* of the axioms. Our axioms specify, for instance, that any possible individual preference is valid. While this is attractive in an axiomatic setting, it is necessary to the Arrow’s Impossibility Result and implies possibly breaking transitivity in light of May’s Theorem. However, those conditions *can* be avoided. In some cases (such as for Corollary 4.13 and Theorem 4.20), a suitable domain restriction, which disallows some particular individual preferences, avoid the issues (please refer to [6, 7]). If preferences are *single-peaked* or *quasi-transitive*, we also dodge from the impossibility results above (please refer to [16, 15]). Given the considerations here, we invite the reader to appreciate these previous results and understand their relevance for algorithm design – not only in the context of protein folding with

the thermodynamic hypothesis, but perhaps in other similar contexts where a 2-step aggregation procedure is similarly suitable. The social choice literature is vast – and we urge the community of computer scientists to investigate these results and their possibly immense connection and relevance to problems involving protein folding.

References

- [1] <http://www.rcsb.org>.
- [2] Christian B. Anfinsen. Principles that govern the folding of protein chains. *Science*, 181(4096):223–230, 1973.
- [3] Kenneth Arrow. *Social Choice and Individual Values*. John Wiley and Sons, 1963.
- [4] K.J. Arrow and M.D. Intriligator, editors. *Handbook of Mathematical Economics*, volume III. Elsevier Science Publishers, 1986.
- [5] Nick Baigent. Preference proximity and anonymous social choice. *The Quarterly Journal of Economics*, 102(1):pp. 161–170, 1987.
- [6] Donald Campbell and Jerry Kelly. Impossibility theorems in the arrowian framework. volume 1 of *Handbook of Social Choice and Welfare*, chapter 1, pages 35 – 94. Elsevier, 2002.
- [7] Graciela Chichilnisky. Social aggregation rules and continuity. *The Quarterly Journal of Economics*, 97(2):pp. 337–352, 1982.
- [8] Claude D’Aspremont and Louis Gevers. Equity and the informational basis of collective choice. *The Review of Economic Studies*, 44(2):pp. 199–209, 1977.
- [9] Ken A. Dill, S. Banu Ozkan, M. Scott Shell, and Thomas R. Weikl. The protein folding problem. *Annual Review of Biophysics*, 37(1):289–316, 2008.
- [10] John Geanakoplos. Three brief proofs of arrow’s impossibility theorem. Cowles Foundation Discussion Papers 1123R3, Cowles Foundation for Research in Economics, Yale University, April 1996.
- [11] R. Jernigan and I. Bahar. Structure-derived potentials and protein simulations. *Current Opinion Structural Biology*, 6(2):195–209, 1996.
- [12] Kenneth May. A set of independent necessary and sufficient conditions for simple majority decision. *Econometrica*, 20(4):680–684, 1952.
- [13] S. Miyazawa and R. Jernigan. Self-consistent estimation of inter-residue protein contact energies based on an equilibrium mixture approximation of residues. *Proteins: Structure, Function, and Genetics*, 34(1):49–68, 1999.
- [14] Baigent Nicholas. Chapter eighteen - topological theories of social choice. In Amartya Sen Kenneth J. Arrow and Kotaro Suzumura, editors, *Handbook of Social Choice and Welfare*, volume 2 of *Handbook of Social Choice and Welfare*, pages 301 – 334. Elsevier, 2011.

- [15] Amartya Sen. Quasi-transitivity, rational choice and collective decisions. *The Review of Economic Studies*, 36(3):pp. 381–393, 1969.
- [16] Amartya K. Sen. A Possibility Theorem on Majority Decisions. *Econometrica*, 34(2):491–499, 1966.
- [17] P. Thomas and K. Dill. An iterative method for extracting energy-like quantities from protein structures. *Proceedings of the National Academy of Sciences of the United States of America*, 93(21):11628–11633, 1996.